

AMENDMENTS TO THE CLAIMS

1. (Currently amended) A method for processing a host command in a storage system including a plurality of storage elements coupled to a plurality of storage controllers, the method comprising the steps of:

at a first one of said plurality of storage controllers,

receiving a host command directed to a volume of said storage system from a host;

storing the host command until the host command has been executed to completion and the execution of the host command has been reported to the host;

determining a target storage element of said storage system corresponding to said volume;

determining a target storage controller from said plurality of storage controllers corresponding to said target storage element;

if said target storage controller is not said first one of said plurality of storage controllers, forwarding the host command to said target storage controller;

at the target storage controller,

receiving the host command forwarded by the first one of said plurality of controllers;

executing said host command;

in response to after said host command being ~~has been~~ executed to completion by the target storage controller, forwarding, to said first one of said plurality of storage controllers, an execution status indicating the completion of said host command; and

at said first one of said plurality of storage controllers,

receiving said execution status;

forwarding said execution status to said host; and

erasing the stored host command,

wherein said step of determining a target storage controller comprises searching a mapping table.

2. (Canceled)

3. (Previously presented) The method of claim 1, further comprising the step of:

before receiving a host command at said first one of said plurality of storage controllers, creating said mapping table to reflect a mapping between volumes and controllers of said storage system.

4. (Currently amended) A method for processing a host command in a storage system including a plurality of storage elements coupled to a plurality of storage controllers, the method comprising the steps of:

at a first one of said plurality of storage controllers,

receiving a host command directed to a virtual volume of said storage system from a host;

storing the host command until the host command has been executed to completion and the execution of the host command has been reported to the host;

determining a plurality of logical volumes of said storage system corresponding to said virtual volume;

determining a plurality of target storage controllers from said plurality of storage controllers corresponding to said plurality of logical volumes;

sending a component command to each of said plurality of target storage controllers, the component command for a particular one of said plurality of target storage controllers corresponding to at least a portion of the host command and relating to the logical volume associated with the particular target controller;

at each of said plurality of the target storage controllers,

receiving the component command sent by the first one of said plurality of controllers;

executing said component command;

in response to ~~after~~ said component command being ~~has been~~ executed to completion by the target storage controller, forwarding, to the first one of said plurality of storage controllers, an execution status indicating the completion of said component command; and

at said first one of said plurality of storage controllers,

receiving said execution status from each one of said plurality of target storage controllers;

determining an aggregate host command execution status from said received execution status;

forwarding said aggregate host command status to said host; and

erasing the stored host command,

wherein said step of determining a plurality of target storage controllers comprises searching a mapping table.

5. (Canceled)

6. (Previously presented) The method of claim 4, further comprising the step of:

before receiving a host command at said first one of said plurality of storage controllers, creating said mapping table to reflect a mapping between virtual volumes and logical volumes of said storage system.

7. (Currently amended) A storage system, comprising:

a plurality of storage controllers, each one of said plurality of storage controllers comprising:

at least one host port for communicating with a plurality of hosts; and

at least one storage element port for communicating with a plurality of storage elements;

an interconnect connected to each of said plurality of storage controllers;

at least one storage element, each storage element being coupled to a respective storage element port; and

a configuration controller for setting up mappings between storage volumes and storage elements,

wherein said plurality of storage controllers further comprise means for processing a host command received from a host on any host port targeted to a storage volume on any one of said plurality of storage elements, the host command being received and stored by a single storage controller of the plurality of storage controllers and an execution status indicating completion of the host command being reported to a host by the single storage controller of the plurality of storage controllers in response to ~~after the host command being has been~~ executed to completion even when the host command is executed by others of the plurality of storage controllers.

8. (Canceled)

9. (Currently amended) A storage system, comprising:

a plurality of storage controllers, each one of said plurality of storage controllers comprising:

at least one host port for communicating with a plurality of hosts; and

at least one storage element ports for communicating with a plurality of storage elements;

an interconnect connected to each of said plurality of storage controllers;

at least one storage element, each storage element being coupled to a respective storage element port; and

a configuration controller for setting up mappings between virtual volumes and logical volumes,

wherein said plurality of storage controllers further comprise means for processing a host command received from a host on any host port targeted to a virtual volume on said plurality of storage elements, the host command being received and stored and an execution status indicating completion of the host command being reported to a host by the single storage controller of the plurality of storage controllers in response to ~~after the host command being~~ ~~has been~~ executed to completion even when the host command is executed by others of the plurality of storage controllers.

10. (Canceled)

11. (Currently amended) A scalable storage controller, comprising:

a high speed interconnect;

a plurality of modules, each of said plurality of modules comprising:

a cache memory;

an interconnect port, coupled to said high speed interconnect;

a storage port, for coupling to a storage device;

a host port, for coupling to a host; and

a processing element, coupled to said cache memory, said interconnect port, said storage port, and said host port,

wherein the processing element of each of said modules is configured to search mappings between storage volumes and storage devices and to receive and store a host command received from a host and report the execution status of the host command to the host in response to ~~after~~ the host command being ~~has been~~ executed to completion even when the host command is executed by others of the plurality of modules, the execution status indicating the completion of the host command.

12. (Original) The scalable storage controller of claim 11, wherein said high speed interconnect is a serial interconnect.

13. (Original) The scalable storage controller of claim 11, wherein said high speed interconnect is a fibre channel loop.

14. (Original) The scalable storage controller of claim 11, wherein said high speed interconnect is a switched serial network.

15. (Currently amended) A scalable storage controller, comprising:

a primary interconnect; and

a plurality of modules, each of said plurality of modules comprising:

a cache memory;

an interconnect port, coupled to said primary interconnect;

a storage port, for coupling to a storage device;

a host port, for coupling to a host; and

a processing element, coupled to said cache memory, said interconnect port, said storage port, and said host port;

wherein the processing element of each of said modules is configured to process a host command received from a host for accessing a virtual volume by causing each module associated with said virtual volume to access its respective storage device, to search mappings between virtual volumes and logical volumes and to store the host command and report the execution status of the host command to the host in response to ~~after~~ the host command being ~~has been~~ executed to completion even when the host command is executed by others of the plurality of modules, the execution status indicating completion of the host command.

16. (Original) The scalable storage controller of claim 15, further comprising:

a redundant interconnect;

wherein each of said modules further comprises another interconnect port coupled to said redundant interconnect.

17. (Currently amended) In a storage controller comprising a plurality of modules each capable of receiving host commands, a method of conflict detection, comprising:

receiving, at a first module, a first access request;

storing the first access request at the first module until the first access request has been completed;

receiving, at a second module, a second access request;

storing the second access request at the second module until the second access request has been completed;

at each of said first and second modules,

determining whether said received access request corresponds to either a read or a write request to two or more storage elements;

if said received access request is determined to correspond to either a read or a write request to two or more storage elements, executing a coordination request, and

if said coordination request is not granted,

indicating in a conflict table that said received access request is in conflict; and

holding said access request;

if said coordination request is granted,

executing said received access request;

searching said conflict table to find an entry in conflict with said received access request;

if said search is successful, notifying the module associated with said found entry;

if said search is unsuccessful, releasing said coordination request; and

in response to after the access request being has been executed to completion, forwarding an execution status from the module that received the access request to a host that sent the access request, the execution status indicating completion of the access request.

18. (Previously presented) The method of claim 17, wherein said step of executing a coordination request comprises:

searching a coordination data structure to find a data structure having an address range encompassing an address range associated with said received access request;

if the data structure is found, attempting to reserve said data structure;

if said attempt to reserve is not successful,

indicating the said coordination request is not granted;

if said attempt to reserve is successful,

indicating that said coordination request is granted;

if said data structure is not found,

allocating said data structure and indicating said coordination request is granted.